

Object-Part Attention Driven Discriminative Localization for Fine-grained Image Classification

Yuxin Peng, Xiangteng He, and Junjie Zhao

Abstract—Fine-grained image classification is to recognize hundreds of subcategories belonging to the same basic-level category, such as 200 subcategories belonging to bird, and highly challenging due to *large* variance in same subcategory and *small* variance among different subcategories. Existing methods generally find where the object or its parts are and then discriminate which subcategory the image belongs to. However, they mainly have *two limitations*: (1) Relying on object or parts annotations which are heavily labor consuming. (2) Ignoring the spatial relationship between the object and its parts as well as among these parts, both of which are significantly helpful for finding discriminative parts. Therefore, this paper proposes the *object-part attention driven discriminative localization (OPADDL)* approach for weakly supervised fine-grained image classification, and the main novelties are: (1) *Object-part attention model* integrates two level attentions: *object-level attention* localizes objects of images, and *part-level attention* selects discriminative parts of object. Both are jointly employed to learn multi-view and multi-scale features to enhance their mutual promotion. (2) *Object-part spatial model* combines two spatial constraints: *object spatial constraint* ensures selected parts highly representative, and *part spatial constraint* eliminates redundancy and enhances discrimination of selected parts. Both are jointly employed to exploit the subtle and local differences for distinguishing the subcategories. Importantly, neither objects nor parts annotations are used, which avoids the heavy labor consuming of labeling. Comparing with more than 10 state-of-the-art methods on 3 widely used datasets, our OPADDL approach achieves the best performance.

Index Terms—Object-part attention model, discriminative localization, object-part spatial constraint model, weakly supervised learning, fine-grained image classification.

I. INTRODUCTION

FINE-GRAINED image classification is highly challenging, aiming to recognize hundreds of subcategories under same basic-level category, such as hundreds of subcategories of birds [1], cars [2], pets [3] and aircrafts [4]. While basic-level image classification only needs to discriminate the basic-level category, such as bird or car. The difference between basic-level and fine-grained image classification is shown as Fig. 1. Fine-grained image classification is a highly important task with wide applications, such as automatic driving, biological conservation and cancer detection. Fig. 2 shows the large variance in same subcategory and small variance among different subcategories, and it is extremely hard even for human to recognize hundreds of subcategories, such as





Inputs				
Basic-level Image Classification	Bird	Bird	Car	Car
Fine-grained Image Classification	American Crow	Fish Crow	Hyundai Elantra Sedan 2007	Toyota Sequoia SUV 2012

Fig. 1. Basic-level image classification vs. fine-grained image classification. In basic-level image classification, it is only needed to classify the first two images to bird category, distinguishing them from car category. While in fine-grained image classification, which subcategory it belongs to should know exactly. For example, the first two images belong to the subcategories of American Crow and Fish Crow respectively.

200 bird subcategories or 196 car subcategories. Due to small variance in the object appearances, subtle and local differences are the key points for fine-grained image classification, such as the color of back, the shape of bill and the texture of feather for bird. Since these subtle and local differences locate at the object and its discriminative parts, most existing methods [5]–[7] generally follow the strategy of finding where the object or its parts are in the image and then discriminating which subcategory the image belongs to.

For finding the object and its discriminative parts, generating image patches with high objectness by a bottom-up process is generally first performed, meaning that the generated patches contain the object or its discriminative parts. Selective search [8] is an unsupervised method that can generate thousands of such image patches, which is extensively used in recent works [5], [6], [9]. Since the bottom-up process has high recall but low precision, it is indispensable to remove the noisy image patches and retain the ones contain object or its discriminative parts through top-down attention model. In the context of fine-grained image classification, finding the object and its discriminative parts can be regarded as a two-level attention process, one is object-level and the other is part-level. An intuitive idea is to use object annotation (i.e. bounding box of object) for object-level attention and parts annotations (i.e. parts locations) for part-level attention. Most existing methods [5], [10]–[12] rely on the object or parts annotations for finding the object or its discriminative parts, but such labeling is heavily labor consuming. This is the *first limitation*.

For addressing the above problem, researchers begin focusing on how to achieve promising performance under the weakly supervised setting that neither object nor parts annotations are used in both training and testing phase. Zhang et al.

This work was supported by National Natural Science Foundation of China under Grants 61371128 and 61532005.

The authors are with the Institute of Computer Science and Technology, Peking University, Beijing 100871, China. Corresponding author: Yuxin Peng (e-mail: pengyuxin@pku.edu.cn).



Fig. 2. Illustration of difficulty in fine-grained image classification: large variance in same subcategory as shown in the first row, and small variance among different subcategories as shown in the second row. The images in (a) Birds, (b) Cars and (c) Cats are from CUB-200-2011 [1], Cars-196 [2] and Oxford-IIIT Pet [3] respectively.

[13] propose to select the discriminative parts through exploiting the useful information in part clusters, and Zhang et al. [6] propose an automatic fine-grained image classification method, incorporating deep convolutional filters for both parts selection and description. However, when they select the discriminative parts, the spatial relationship between the object and its parts as well as among these parts are ignored, but both of them are significantly helpful for finding the discriminative parts. This causes the selected parts: (1) have large areas of background noise and small areas of object, (2) have large overlap with each other which leads to the redundant information. This is the *second limitation*.

For addressing the above two limitations, this paper proposes the object-part attention driven discriminative localization (OPADDL) approach for weakly supervised fine-grained image classification. Its main novelties and contributions can be summarized as follows:

- **Object-Part Attention Model.** Most exiting works rely on object or parts annotations [5], [11], [12], [14], while labeling is heavily labor consuming. For addressing this important problem, we propose the object-part attention model for weakly supervised fine-grained image classification to avoid using the object or parts annotations and march toward practical applications. It integrates two level attentions: (1) **Object-level attention model** utilizes the global average pooling in CNN to extract the saliency map for localizing objects of images, which is to learn object features. (2) **Part-level attention model** first selects the discriminative parts and then aligns the parts based on the cluster pattern of neural network, which is to learn the subtle and local features. The object-level attention model focuses on the representative object appearance, and the part-level attention model focuses on the distinguishing specific differences of parts among subcategories. Both of them are jointly employed to boost the multi-view and multi-scale feature learning, and enhance their mutual promotion to achieve the good performance for fine-grained image classification.
- **Object-Part Spatial Constraint Model.** Most existing weakly supervised methods [6], [13] ignore the spatial

relationship between the object and its parts as well as among these parts, both of which are significantly helpful for discriminative parts selection. For addressing this problem, we propose the parts selection approach driven by object-part spatial model, which combines two types of spatial constraints: (1) **Object spatial constraint** enforces that the selected parts are located in the object region to ensure their high representativeness. (2) **Part spatial constraint** reduces the overlap among parts and highlights the saliency of parts to eliminate the redundancy and enhance the discrimination of selected parts. Combining these two spatial constraints not only exploits the subtle and local discrimination for promoting parts selection significantly, but also achieves a notable improvement on fine-grained image classification.

Our previous conference paper [15] integrates two level attentions: object-level attention selects image patches relevant to the object and part-level attention selects discriminative parts, which is the first work to classify fine-grained images without using object and parts annotations in both training and testing phase, and achieves promising results [13]. In this paper, our OPADDL approach further exploits the two level attentions for discriminative localization, exploiting not only the discriminative parts but also the object, and employing the object-part spatial constraint model to eliminate redundancy as well as highlight discrimination of the selected parts: **For object-level attention**, we further propose an *automatic object localization approach via saliency extraction* to focus on the representative object feature for better classification performance. It utilizes the global average pooling in CNN for localizing objects of images, rather than only selecting the image patches relevant to object that have large areas of background noise or not contain the whole object in image like [15]. **For part-level attention**, we further propose a *parts selection approach driven by object-part spatial constraint model* to exploit the subtle and local differences among subcategories. It considers the spatial relationship between object and its parts as well as among these parts, rather than ignoring both two spatial constraints that generates large areas of background noise and large overlap among selected

parts like [15]. Compared with more than 10 state-of-the-art methods on 3 widely used datasets, the effectiveness of our OPADDL approach is shown from the comprehensive experimental results.

The rest of this paper is organized as follows: Section II briefly reviews related works on fine-grained image classification. Section III presents our proposed OPADDL approach, and Section IV introduces the experiments as well as the results analyses. Finally Section V concludes this paper.

II. RELATED WORK

Most traditional methods for fine-grained image classification follow the strategy of extracting basic low-level descriptors like SIFT [16], and then generating Bag-of-Words for image representation [17], [18]. However, the performance of these methods is limited by the handcrafted features. Deep learning has shown its strong power in feature learning, and achieved great progresses in fine-grained image classification [5]–[7], [10], [15], [19]–[25]. These methods can be divided into 3 groups [26]: (1) ensemble of networks based methods; (2) visual attention based methods; (3) part detection based methods.

A. Ensemble of Networks Based Methods

Ensemble of networks based methods are proposed to utilize multiple neural networks to learn different representations of image for better classification performance. Each subcategory has an implied hierarchy of labels in its ontology tree. For example, *Picoides Pubescens*, which is the label of specie level, has the label of genus level as *Picoides* and the family level as *Picidae*. Wang et al. [24] first leverages the labels of multiple levels to train a series of CNNs at each level, which focuses on different regions of interest in images. Different features are extracted by different level CNNs, and combined to encode informative and discriminative features. Finally, a linear SVM is trained to learn weights for the final classification. However, the external labels of ontology tree are needed to know. Lin et al. [25] propose a bilinear CNN model, which is an end-to-end system jointly combining two CNNs, each of which is adopted as a feature extractor. The extracted features at each location of image are multiplied by outer product processing, and then pooled to generate an image descriptor. Finally, softmax is conducted for final prediction. Despite achieving promising results, these methods are still limited by the lack of ability to be spatially invariant to the input image. Therefore, Jaderberg et al. [21] propose a learnable network, called spatial transformer, which consists of three parts: localization network, grid generator and sampler. Four spatial transformers in parallel are performed on images, and capture the discriminative parts to pass to the part description subnets. Finally, softmax is conducted on the concatenated part descriptor for final prediction.

B. Visual Attention Based Methods

Due to attention system, humans focus on the discriminative regions of an image dynamically, rather than receive and deal

with the information of entire image directly. The natural advantage makes the attention mechanism widely used in fine-grained image classification. Inspired by the way how humans perform visual sequence recognition, Sermanet et al. [27] propose the attention for fine-grained categorization (AFGC) system. First, process a multi-resolution crop on the input image, each crop is called glimpse. And then use the information of glimpse to output the next location and possibly the next object via a deep recurrent neural network at each step. The final prediction is computed through the sequence of glimpses. Recently, fully convolutional neural network is used to learn the saliency of an image for finding the discriminative regions [28]. Liu et al. [29] uses the fully convolutional attention to localize multiple object parts to get better classification performance. Zhou et al [28] uses global average pooling (GAP) in CNN to generate the saliency map for each image. Based on the saliency map, the discriminative region can be found. And furthermore, a diversified visual attention network (DVAN) [30] is proposed to pursue the diversity of attention as well as gather discriminative information. In this paper, our OPADDL approach integrates two level attention models: object-level attention model focuses on the representative object appearance, and part-level attention model focuses on the discriminative parts. Both of them are jointly employed to learn multi-view and multi-scale features to enhance their mutual promotion.

C. Part Detection Based Methods

In fine-grained image classification, subtle and local differences are generally located at its discriminative parts of object, so the discriminative part detection is crucial for fine-grained image classification. Girshick et al. [9] propose the popular detection method, R-CNN. First generate thousands of candidate image patches for each image via the bottom-up process [8], and then select the image patches with high classification scores as detection results. Zhang et al. [5] utilize R-CNN with a geometric prior to detect discriminative parts in fine-grained image classification, and then train a classifier on the features of detected parts for final categorization. They use both the object and parts annotations. Recently, researchers begin focusing on how to detect the discriminative parts under the weakly supervised setting, which means neither object nor parts annotations are used in both training and testing phase. Simon and Rodner [20] propose a constellation model to localize parts of object, leveraging CNN to find the constellations of neural activation patterns. First, neural activation maps are computed as a part detector by using the outputs of a middle layer of CNN. Second, a part model is estimated by selecting part detectors via constellation model. Finally, the part model is used to extract features for classification. Zhang et al. [6] propose an automatic fine-grained image classification method, incorporating deep convolutional filters for both parts selection and description. They combine two steps of deep filter response picking: The first step picks the discriminative filters that significantly respond to specific parts in image. The second step picks the salient regions and generates features with spatially weighted Fisher Vector based on the saliency

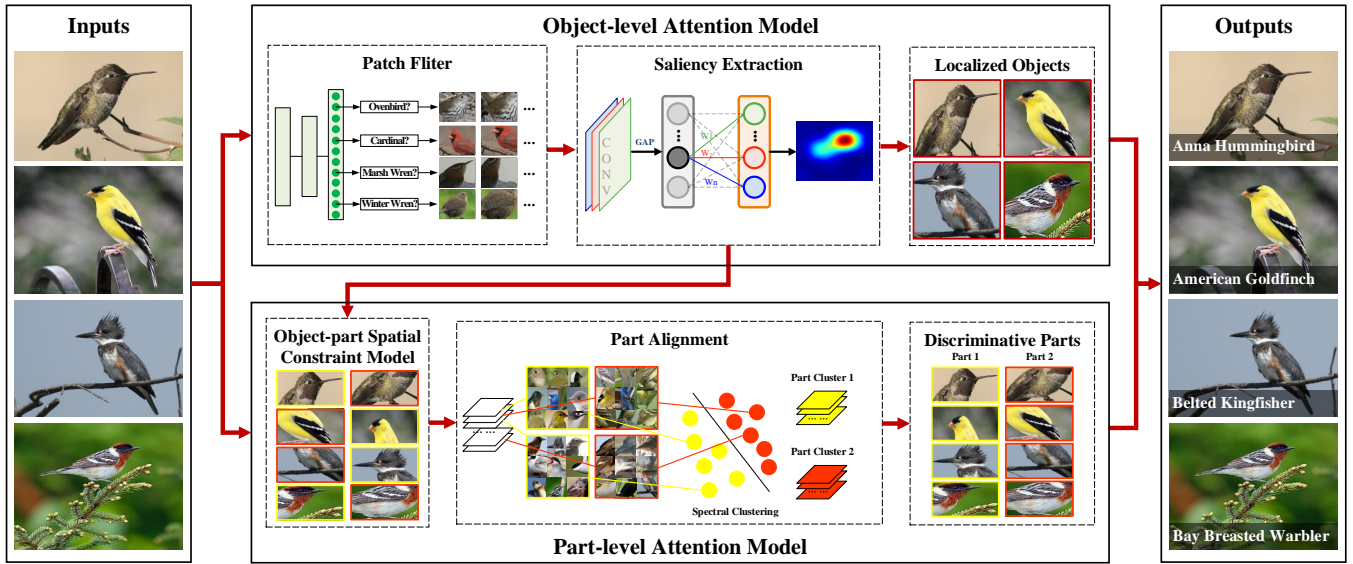


Fig. 3. An overview of our OPADDL approach. The object-level attention model is to localize object for learning object features. The part-level attention model is to select the discriminative parts for exploiting the subtle and local features. The outputs show the predicted subcategories.

map for classification. Zhang et al. [13] propose to select the discriminative parts through exploiting the useful information in part clusters. In our OPADDL approach, we first propose an object-part spatial constraint model to select discriminative parts, which considers the spatial relationship between object and its parts as well as among these parts, and then utilize the cluster pattern of neural network to align the parts with same semantic meaning together for improving the classification performance.

III. OUR OPADDL APPROACH

Our approach is based on an intuitive idea: fine-grained image classification generally first localizes the object (object-level attention) and then its discriminative parts (part-level attention). For example, recognizing an image whether contains a Field Sparrow follows the processes of first finding a bird, and then focusing on its discriminative parts that distinguish it from other bird subcategories. We propose the *object-part attention driven discriminative localization* approach for weakly supervised fine-grained image classification, which uses neither object nor parts annotations in both training and testing phase, and only uses the image-level subcategory labels. As shown in Fig. 3, our OPADDL approach first localizes objects of images through object-level attention model for learning object features, and then selects the discriminative parts of object through part-level attention model for learning the subtle and local features. In the following subsections, the object-level and part-level attention models are presented respectively.

A. Object-level Attention Model

Most existing weakly supervised works [6], [13], [20] devote to the discriminative parts selection, but ignore the

object localization, which can remove the influence of background noise in image to learn meaningful and representative object features. Although some methods consider both object localization and parts selection, they rely on the object and parts annotations [5], [19]. For addressing this important problem, we propose an object-level attention model based on the saliency extraction for localizing the objects of images automatically only with image-level subcategory labels, without using any object or parts annotations. The model consists of two components: patch filter and saliency extraction. The first component is to filter out the noisy image patches and retain these ones relevant to the object for training a CNN called *ClassNet*, to learn multi-view and multi-scale features for the specific subcategory. Then the second component is to extract the saliency map via global average pooling in CNN for localizing the object of image.

1) *Patch Filter*: The number of training data is highly helpful for the performance of CNN, so we first focus on how to extend the training data. The bottom-up process can generate thousands of candidate image patches by grouping pixels into regions that may contain the object. These image patches can be used as the extension of training data due to their relevances to the object. Therefore, selective search [8] is adopted to generate candidate image patches for a given image, which is an unsupervised and widely used bottom-up process method. These candidate image patches provide multiple views and scales of original image, which benefit for training a effective CNN to achieve better fine-grained image classification. However, these patches can not be directly used due to the high recall but low precision, which means a large number of noises exist. The object-level attention model is highly helpful for selecting the patches relevant to the object.

We remove the noisy patches and select relevant patches through a CNN, called *FilterNet*, which is pre-trained on the

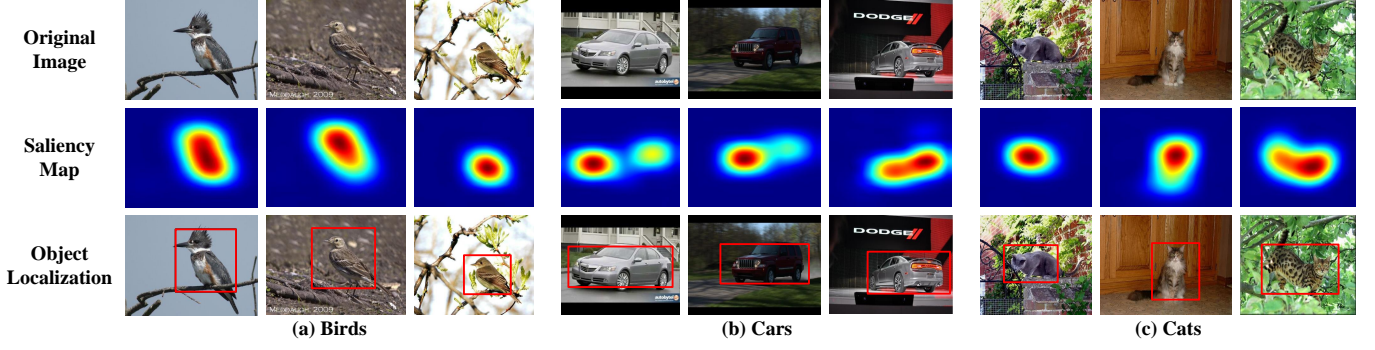


Fig. 4. Some results of saliency extraction by our OPADDL approach. The first row shows the original images and the second row shows the saliency maps of original images. The object localization results are shown in the third row, in which the red rectangles represent the bounding boxes automatically produced by saliency extraction. The images in (a) Birds, (b) Cars and (c) Cats are from CUB-200-2011 [1], Cars-196 [2] and Oxford-IIIT Pet [3] respectively.

ImageNet 1K dataset [31], and then fine-tuned on the training data. We define the activation of neuron in softmax layer belonging to the subcategory of input image as the selection confidence score, and then a threshold is set to decide whether a given candidate image patch should be selected or not.

Through this object-level attention selection, we obtain the image patches relevant to the object with multiple views and multiple scales. The extension of training data improves the training effect of *ClassNet*, which has two aspects of benefits for our OPADDL approach: (1) *ClassNet* is an effective fine-grained image classifier itself, and (2) its internal features are significantly helpful to build part clusters for aligning the parts with same semantic meaning together, which will be described latter in Subsection B. It is noted that the patch filter is performed only in the training phase and only uses image-level subcategory label.

2) *Saliency Extraction*: In this stage, CAM [28] is adopted to obtain the saliency map M_c of an image for subcategory c to localize the object. The saliency map indicates the representative regions used by the CNN to identify the subcategory of image, as shown in the second row of Fig. 4. Then object regions of images, as shown in the third row of Fig. 4, can be obtained by performing binarization and connected region extraction on the saliency maps.

Given an image I , the activation of neuron u in the last convolutional layer at spatial location (x, y) is defined as $f_u(x, y)$, and w_u^c defines the weight corresponding to subcategory c for neuron u . The saliency value at spatial location (x, y) for subcategory c is computed as follows:

$$M_c(x, y) = \sum_u w_u^c f_u(x, y) \quad (1)$$

where $M_c(x, y)$ directly indicates the importance of activation at spatial location (x, y) leading to the classification of an image to subcategory c . Instead of using the image-level subcategory labels, we use the prediction result as the subcategory c used in saliency extraction for each image.

Through object-level attention model, we localize objects in images to train a CNN called *ObjectNet* for obtaining the prediction of object-level attention.

B. Part-level Attention Model

Since the discriminative parts, such as head and body, are crucial for fine-grained image classification, previous works [5], [10], [19] select discriminative parts from the candidate image patches produced by the bottom-up process like selective search [8]. However, these works rely on the parts annotations which are heavily labor consuming. Although some works begin to focus on finding the discriminative parts without using any parts annotations [6], [15], they ignore the spatial relationship between the object and its parts as well as among these parts. Therefore, we propose a new parts selection approach driven by object-part spatial constraint model for exploiting the subtle and local discrimination for distinguishing the subcategories, which uses neither object nor parts annotations. It consists of two components: object-part spatial constraint model and part alignment. The first is to select the discriminative parts, and the second is to align the selected parts into clusters by the semantic meaning.

1) *Object-Part Spatial Constraint Model*: We obtain object regions of images through object-level attention model, and then employ object-part spatial constraint model to drive the parts selection from the candidate image patches produced by the bottom-up process. Two spatial constraints are jointly considered: the spatial relationship between object and its parts called *object spatial constraint*, and the spatial relationship among these parts called *part spatial constraint*. For a given image I , its saliency map M and object region b are obtained through object-level attention model. And then parts selection is driven by object-part spatial constraint model as follows:

Let \mathbb{P} denotes all the candidate image patches and $P = \{p_1, p_2, \dots, p_n\}$ denotes the n parts we selected from \mathbb{P} as the discriminative parts for each given image. The object-part spatial constraint model considers the combination of two spatial constraints by solving the following optimization problem:

$$P^* = \arg \max_{\mathbb{P}} \Delta(P) \quad (2)$$

where $\Delta(P)$ is defined as a scoring function over two spatial constraints as follows:

$$\Delta(P) = \Delta_{box}(P) \Delta_{parts}(P) \quad (3)$$

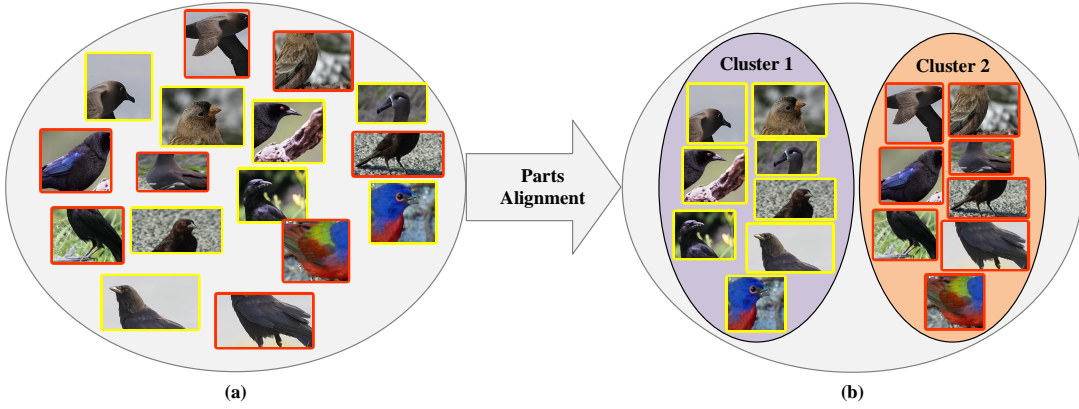


Fig. 5. Some results of part alignment in our OPADDL approach. (a) shows the image patches which are selected through object-part spatial constraint model, and (b) shows that the image patches are aligned into clusters via part clusters.

where $\Delta_{box}(P)$ denotes the object spatial constraint and $\Delta_{parts}(P)$ denotes the part spatial constraint, described in the following paragraphs.

Object spatial constraint. Ignoring the spatial relationship between the object and its parts causes that the selected parts have large areas of background noise but small areas of part, which decreases the representativeness of selected parts. Since the discriminative parts are inside the object region, an intuitive spatial constraint function is defined as:

$$\Delta_{box}(P) = \prod_{i=1}^n f_b(p_i) \quad (4)$$

where

$$f_b(p_i) = \begin{cases} 1 & \text{IoU}(p_i) > \text{threshold} \\ 0 & \text{otherwise} \end{cases} \quad (5)$$

and $\text{IoU}(p_i)$ defines the proportion of Intersection-over-Union (IoU) overlap of part region and object region. It is noted that the object region is obtained automatically through the object-level attention model, *not provided by the object annotation*. Neither object nor parts annotations are used in any phase of our approach.

Part spatial constraint. Ignoring the spatial relationship among these parts leads to the problem that the selected parts have large overlap with each other and some discriminative parts are ignored. The saliency map indicates the discrimination of image and benefits for selecting discriminative parts. We jointly model saliency and the spatial relationship of parts as follows:

$$\Delta_{parts}(P) = \log(A_U - A_I - A_O) + \log(\text{Mean}(M_{A_U})) \quad (6)$$

where A_U is the union area of n parts, A_I is the intersection area of n parts, A_O is the area outside the object region and $\text{Mean}(M_{A_U})$ is defined as follows:

$$\text{Mean}(M_{A_U}) = \frac{1}{|A_U|} \sum_{i,j} M_{ij} \quad (7)$$




Fig. 6. Illustration of spectral clustering.

where pixel (i, j) locates in the union area of parts, M_{ij} refers to the saliency value of pixel (i, j) , and $|A_U|$ refers to the number of pixels that locate in the union area of n parts.

2) Part Alignment: The selected parts through object-part spatial constraint model are in disorder and not aligned by its semantic meaning, as shown in Fig. 5(a). These parts with different semantic meanings contribute the final prediction differently, so an intuitive idea is to align the parts with the same semantic meaning together, as shown in Fig. 5(b).

We are inspired by the fact that middle layers of *ClassNet* have shown clustering patterns. For example, there are groups of neurons significantly responding to the head of bird, and others to the body of bird, despite the fact they may correspond to different poses. So clustering is performed on the neurons of a middle layer in the *ClassNet* to build the part clusters for aligning the selected parts. First compute the similarity matrix S , where $S(i, j)$ denotes the cosine similarity of weights of two mid-layer neurons u_i and u_j . And then perform spectral clustering on the similarity matrix S to partition the mid-layer neurons into m groups. In the experiments, neurons are picked from the penultimate convolutional layer with m set as 2, as

shown in Fig. 6, where the coordinate values represent the two largest eigenvectors of similarity matrices among all neurons.

The processes of using the part clusters to align the selected parts are: (1) Warping the images of selected parts to the size of receptive field on input image of neuron in penultimate convolutional layer. (2) Feed-forwarding the selected parts to the penultimate convolutional layer to produce an activation score for each neuron. (3) Summing up the scores of neurons in one cluster to get cluster score. (4) Aligning the selected parts to the cluster with highest cluster score, and finally the selected parts are aligned into clusters with certain semantic meanings, which is formulated as the following: For a given image, n image patches $P = \{p_1, p_2, \dots, p_n\}$ are obtained by object-part spatial constraint model, and then part alignment is performed on these patches with m part clusters $L = \{l_1, l_2, \dots, l_m\}$ as Algorithm 1.

Algorithm 1 Part Alignment

Input: The i th selected part p_i ; The part clusters $L = \{l_1, l_2, \dots, l_m\}$; And the number of neurons in penultimate convolutional layer d .

Output: The cluster that p_i is aligned into l_c .

```

1: Set  $score_k = 0$ ;  $k = 1, \dots, m$ .
2: Warp  $p_i$  to the size of receptive field on input image of neuron in penultimate convolutional layer.
3: Perform a feed-forward pass to compute  $p_i$ 's activations  $F_i = \{f_{i1}, f_{i2}, \dots, f_{id}\}$ .
4: for  $k = 1, \dots, m$ ;  $j = 1, \dots, d$  do
5:   if  $j$ th neuron belongs to cluster  $l_k$  then
6:      $score_k = score_k + f_{ij}$ .
7:   end if
8: end for
9:  $c = \arg \max_k score_k$ .
10: return  $l_c$ .
```

The choice of neuron layer has an influence on the cluster and classification performance. We follow standard practice and withhold a validation set of 10% training data for grid search to determine which layer to be chosen. At last, we find the penultimate convolutional layer works better than others.

Through part-level attention model, we select the discriminative parts in images to train a CNN called *PartNet* for obtaining the prediction of part-level attention.

C. Final Prediction

For better classification performance, we fine-tune *ClassNet* with the localized object and its discriminative parts to get two classifier, called *ObjectNet* and *PartNet* respectively. *ClassNet*, *ObjectNet* and *PartNet* are all fine-grained image classifiers: *ClassNet* for original images, *ObjectNet* for objects and *PartNet* for selected discriminative parts. However, their functionalities and strengths are different, primarily because they focuses on the different natures of image.

Object-level attention model first drives *FilterNet* to select image patches with multiple views and scales that are relevant to the object, as shown in Fig. 7 (a). These image patches drive *ClassNet* to learn more representative features and localize

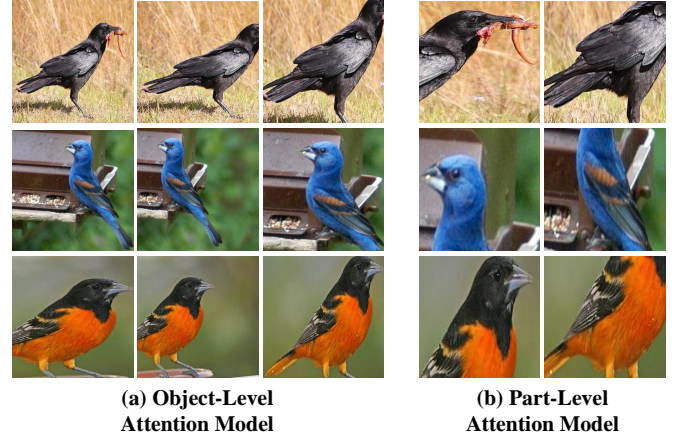


Fig. 7. Some results of selected image patches by the object-level and part-level attention model respectively. The image patches selected by the object-level attention model contain objects, as shown in (a). And the patches selected by the part-level attention model which contain subtle and local features, as shown in (b).

the object region through saliency extraction. And part-level attention model selects discriminative parts which contain subtle and local features, as shown in Fig. 7 (b). The different level focuses (i.e. original image, object of original image, and parts of original image) have different representations and can potentially improve the prediction. Finally, we merge the prediction results of the three different level focuses by using the following equation:

$$final_score = \alpha * original_score + \beta * object_score + \gamma * part_score \quad (8)$$

where $original_score$, $object_score$ and $part_score$ are the softmax values of *ClassNet*, *ObjectNet* and *PartsNet* respectively, and α , β and γ are selected by using the cross-validation method. The subcategory with the highest $final_score$ is chosen as the final prediction result.

IV. EXPERIMENTS

We conduct experiments on 3 widely used datasets for fine-grained image classification: CUB-200-2011, Cars-196 and Oxford-IIIT Pet. Our proposed OPADDL approach is compared with more than 10 state-of-the-art methods to verify its effectiveness.

A. Datasets and evaluation metric

Three datasets are adopted for the experiments:

- **CUB-200-2011** [1]: It is the most widely used dataset for fine-grained image classification, and contains 11788 images of 200 different bird subcategories, which is divided as follows: 5994 images for training and 5794 images for testing. For each subcategory, 30 images are selected for training and 11~30 images for testing. And each image has detailed annotations: a subcategory label, a bounding box of object, 15 part locations and 312 binary attributes. All attributes are visual in nature, pertaining to color, pattern, or shape of a particular part.

- **Cars-196** [2]: It contains 16185 images of 196 car subcategories, and is divided as follows: 8144 images for training and 8041 images for testing. For each subcategory, 24~84 images are selected for training and 24~83 images for testing. And each image is annotated with a subcategory label and a bounding box of object.
- **Oxford-IIIT Pet** [3]: It is a collection of 7349 images with 37 different pet subcategories, among which 12 are cat subcategories and 25 are dog categories. It is divided as follows: 3680 images for training and 3669 images for testing. For each subcategory, 93~100 images are selected for training and 88~100 images for testing. Each image is annotated with a class label, a pixel level segmentation marking the body and a tight bounding box of head.

Accuracy is adopted as the evaluation metric to comprehensively evaluate the classification performances of our OPADDL approach and compared methods, which is widely used for evaluating the performance of fine-grained image classification [5], [6], [13], and is defined as follows:

$$Accuracy = \frac{R_a}{R} \quad (9)$$

where R means the number of testing images and R_a counts the number of images which are correctly classified.

B. Details of the networks

In the experiments, the widely used CNN of VGGNet [32] is adopted. It is noted that the CNN used in our proposed approach can be replaced with the other CNN. In our approach, CNN has two different effects: localization and classification. Therefore, the architectures of CNNs are modified for different functions:

1) *Localization*: In the object-level attention model, CNN is used to extract the saliency map of an image for object localization. Zhou et al. [28] find that the accuracy of localization will be improved if the last convolutional layer before global average pooling has a higher spatial resolution, which is termed as the mapping resolution. In order to get a higher spatial resolution, the layers after conv5-3 are removed, resulting in a mapping resolution of 14×14 . Besides, a convolutional layer of size 3×3 , stride 1, pad 1 with 1024 neurons is added, followed by a global average pooling layer and a softmax layer. The modified VGGNet is pre-trained on the 1.3M training images of ImageNet 1K dataset [31], and then fine-tuned on the fine-grained image classification dataset. The number of neurons in softmax layer is set as the number of subcategories.

2) *Classification*: The CNN used in the experiments for classification is the VGGNet [32] with batch normalization [33]. For the prediction results of original image, object and parts, the same CNN architecture is used but fine-tuned on different training data. For the prediction of original image, we fine-tune the CNN on the image patches selected through object-level attention model, as *ClassNet*. And for the predictions of object and its parts, we fine-tuned the CNNs on the images of objects and images of parts based on *ClassNet* respectively, called *ObjectNet* and *PartNet*.

C. Comparisons with the state-of-the-art methods

This subsection presents the experimental results and analyses of our OPADDL approach on 3 widely used fine-grained image classification datasets as well as the state-of-the-art methods. Table I shows the comparison results on CUB-200-2011. The object, parts annotations and CNN used in these methods are listed for fair comparison. Early works [14], [34], [35] choose SIFT [16] as features, and the performances are limited and much lower than our OPADDL approach no matter whether using the object and parts annotations or not. Our approach is the best among all methods under the same setting that neither object nor parts annotations are used in both training and testing phase, and obtains 1.20% higher accuracy than the best result of FOAF [7] (85.83% vs. 84.63%). It is noted that the CNN used in FOAF is pre-trained not only on ImageNet 1K dataset [31] but also on the dataset of PASCAL VOC [36], while our approach does not use the external dataset like PASCAL VOC. Compared with the second highest result of PD [6], our approach achieves 1.29% higher accuracy (85.83% vs. 84.54%). Our OPADDL approach improves 7.93% than our previous conference paper [15], and it proves the effectiveness of further exploitation in our OPADDL approach, which jointly integrates the object-level and part-level attention models to boost the multi-view and multi-scale feature learning and enhance their complementarity. Besides, our OPADDL approach employs the object-part spatial constraint model to exploit the subtle and local discrimination for distinguishing the subcategories.

And our approach performs better than the methods which focus on the CNN architectures, such as STN [21] and Bilinear-CNN [25]. In STN, GoogleNet [37] with batch normalization [33] is adopted to achieve the accuracy of 82.30% only fine-tuned on the CUB-200-2011 without any other processing. Two different CNNs are employed in Bilinear-CNN: VGGNet [32] and VGG-M [38]. The classification accuracies of the two methods are both 84.10%, which is lower than our approach by 1.73%.

Furthermore, our approach even outperforms the methods which use object annotation, such as Coarse-to-Fine [39], PG Alignment [11] and VGG-BGLM [12]. Moreover, our approach outperforms methods that use both object and parts annotations [5], [23]. Neither object nor parts annotations are used in our OPADDL approach, which lead fine-grained image classification to march toward practical application.

Besides, the results on Cars-196 and Oxford-IIIT Pet are shown in Tables II and III respectively. The trends of results on these two datasets are similar as CUB-200-2011 dataset, our OPADDL approach achieves the best results among state-of-the-art methods (92.19% and 93.81% respectively) and brings 0.89% and 2.21% improvements than the best results of compared methods respectively.

D. Performances of components in our OPADDL approach

Detailed experiments are performed on our OPADDL approach from the following three aspects:

TABLE I
COMPARISONS WITH STATE-OF-THE-ART METHODS ON CUB-200-2011.

Method	Train Annotation		Test Annotation		Accuracy (%)	Net
	Object	Parts	Object	Parts		
Our OPADDL Approach					85.83	VGGNet
FOAF [7]					84.63	VGGNet
PD [6]					84.54	VGGNet
STN [21]					84.10	GoogleNet
Bilinear-CNN [25]					84.10	VGGNet&VGG-M
Multi-grained [24]					81.70	VGGNet
NAC [20]					81.01	VGGNet
PIR [13]					79.34	VGGNet
TL Atten [15]					77.90	VGGNet
MIL [40]					77.40	VGGNet
VGG-BGLm [12]					75.90	VGGNet
Dense Graph Mining [35]					60.19	
Coarse-to-Fine [39]	✓				82.50	VGGNet
Coarse-to-Fine [39]	✓		✓		82.90	VGGNet
PG Alignment [11]	✓		✓		82.80	VGGNet
VGG-BGLm [12]	✓		✓		80.40	VGGNet
Triplet-A (64) [41]	✓		✓		80.70	GoogleNet
Triplet-M (64) [41]	✓		✓		79.30	GoogleNet
Webly-supervised [42]	✓	✓			78.60	AlexNet
PN-CNN [10]	✓	✓			75.70	AlexNet
Part-based R-CNN [5]	✓	✓			73.50	AlexNet
SPDA-CNN [23]	✓	✓	✓		85.14	VGGNet
Deep LAC [43]	✓	✓	✓		84.10	AlexNet
SPDA-CNN [23]	✓	✓	✓		81.01	AlexNet
PS-CNN [22]	✓	✓	✓		76.20	AlexNet
PN-CNN [10]	✓	✓	✓	✓	85.40	AlexNet
Part-based R-CNN [5]	✓	✓	✓	✓	76.37	AlexNet
POOF [34]	✓	✓	✓	✓	73.30	
GPP [14]	✓	✓	✓	✓	66.35	

TABLE II
COMPARISONS WITH STATE-OF-THE-ART METHODS ON CARS-196.

Method	Train Annotation		Test Annotation		Accuracy (%)	Net
	Object	Parts	Object	Parts		
Our OPADDL Approach					92.19	VGGNet
Bilinear-CNN [25]					91.30	VGGNet&VGG-M
DVAN [30]					87.10	VGGNet
FT-HAR-CNN [44]					86.30	AlexNet
HAR-CNN [44]					80.80	AlexNet
PG Alignment [11]	✓				92.60	VGGNet
ELLF [45]	✓				73.90	CNN
R-CNN [9]	✓				57.40	AlexNet
PG Alignment [11]	✓		✓		92.80	VGGNet
BoT(CNN With Geo) [46]	✓		✓		92.50	VGGNet
DPL-CNN [47]	✓		✓		92.30	VGGNet
VGG-BGLm [12]	✓		✓		90.50	VGGNet
LLC [48]	✓		✓		69.50	
BB-3D-G [2]	✓		✓		67.60	

TABLE III
COMPARISONS WITH STATE-OF-THE-ART METHODS ON OXFORD-IIIT PET.

Method	Accuracy (%)	Net
Our OPADDL Approach	93.81	VGGNet
NAC [20]	91.60	VGGNet
FOAF [7]	91.39	VGGNet
ONE+SVM [49]	90.03	VGGNet
Deep Optimized [50]	88.10	AlexNet
NAC [20]	85.20	AlexNet
MsML+ [51]	81.18	CNN
MsML [51]	80.45	CNN
Deep Standard [50]	78.50	AlexNet
Shape+Appearance [3]	56.68	
Zernike+SCC [52]	59.50	
GMP+p [53]	56.80	
GMP [53]	56.10	
M-HMP [54]	53.40	
Detection+Segmentation [55]	54.30	

TABLE IV
PERFORMANCES OF COMPONENTS IN OUR OPADDL APPROACH ON CUB-200-2011, CARS-196 AND OXFORD-IIIT PET.

Method	Accuracy (%)		
	CUB-200-2011	Cars-196	Oxford-IIIT Pet
Our OPADDL Approach (Original+Object-level+Part-level)	85.83	92.19	93.81
Original	80.82	86.79	88.14
Object-level	83.74	88.79	90.98
Part-level	80.65	84.26	85.75
Original+Object-level	84.79	91.15	92.20
Original+Part-level	84.41	91.06	91.82
Object-level+Part-level	84.73	89.69	91.50

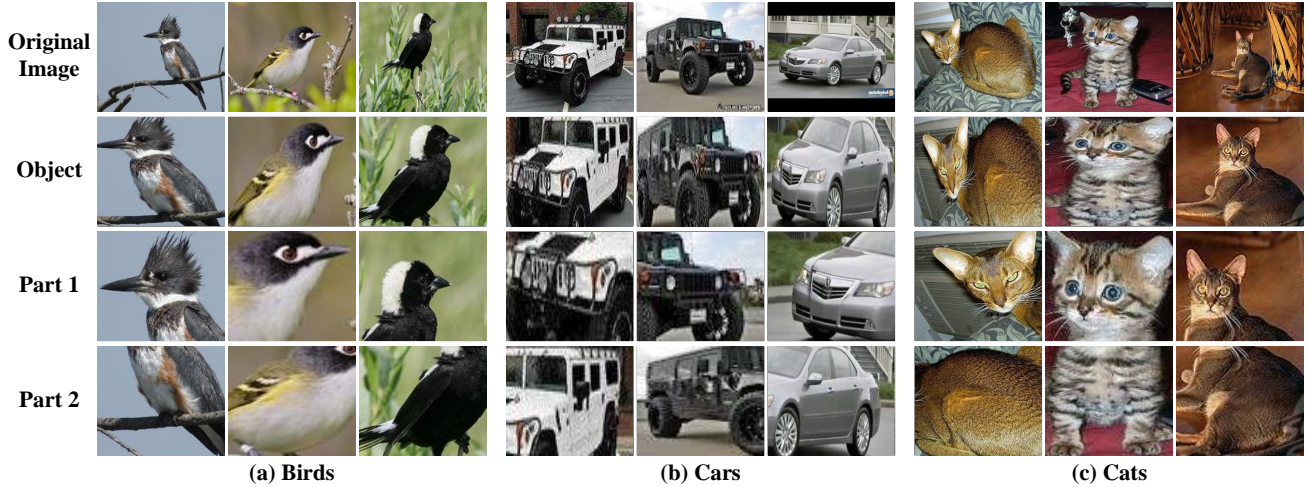


Fig. 8. Some results of object localization and parts selection. The first row denotes the original images, the second row denotes the localized objects of original images via object-level attention model, and the third and fourth rows denote the selected discriminative parts via part-level attention model. The images in (a) Birds, (b) Cars and (c) Cats are from CUB-200-2011 [1], Cars-196 [2] and Oxford-IIIT Pet [3] respectively.

1) *Effectivenesses of object-level attention and part-level attention models:* In our OPADDL approach, the final prediction score is generated by merging the prediction scores of three different images, i.e. original image, image of object and images of parts, which are denoted as “Original”, “Object-level” and “Part-level”. The effectivenesses of object-level and part-level attention models are verified in the following paragraphs. From table IV and Fig. 8, we can observe that:

- Object-level attention model improves the classification accuracy via localizing objects of images for learning global features. Comparing with the result of “Original”, it improves 2.92%, 2.00% and 2.84% on 3 datasets respectively, and combining “Object-level” with “Original” improves even more, i.e. 3.97%, 4.36% and 4.06% on 3 datasets respectively. The classification accuracy of part-level attention model is not higher than “Original”, but combining “Part-level” with “Original” brings 3.59%, 4.27% and 3.68% improvements respectively. The improvements are due to the subtle and local features learned by our part-level attention model.
- Combining object-level and part-level attention models achieves more accurate results than only one level attention model is adopted, e.g. 84.73% vs. 83.74% and 80.65% on the CUB-200-2011 dataset. And combining the two level attention models with “Original” improves

a lot than “Original”, i.e. 5.01%, 5.40% and 5.67% on the 3 datasets respectively. It shows the complementarity of object-level and part-level attention models in fine-grained image classification. The two level attention models have different but complementary focuses: the object-level attention model focuses on differences of representative object appearances, while the part-level attention model focuses on the subtle and local differences of discriminative parts among subcategories. Both of them are jointly employed to boost the multi-view and multi-scale feature learning and enhance their mutual promotion to achieve the better performance for fine-grained image classification.

- Fig. 8 shows some results of object localization and parts selection by our OPADDL approach. The first row denotes the original images, the second row denotes the localized objects of original images via object-level attention model, and the third and fourth rows denote the selected discriminative parts via part-level attention model. The selected parts have explicit semantic meanings, where the third row denotes the head of object and the fourth denotes the body. It is noted that neither object nor parts annotations are used in our OPADDL approach, which avoid the heavy labor consuming of labeling as well as push fine-grained image classification towards practical applications.

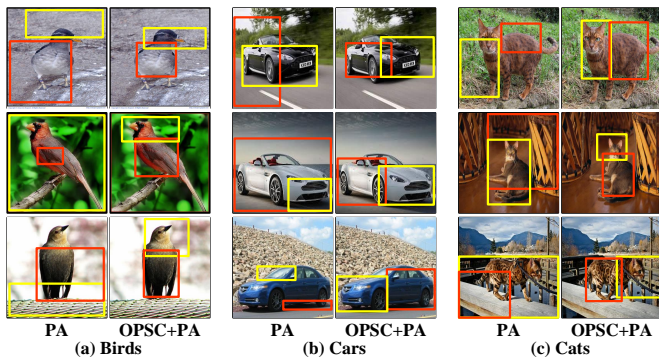


Fig. 9. Examples of parts selection from our previous conference paper [15] (left column) and our OPADDL approach in this paper (right column). “PA” refers to part alignment which is adopted in our previous conference paper [15], “OPSC” refers to object-part spatial constraint model, and “OPSC+PA” refers to combining the above two approaches, which is adopted in our OPADDL approach. The yellow and orange rectangles denote the selected discriminative parts via the two approach, which respond to the heads and bodies of objects. The images in (a) Birds, (b) Cars and (c) Cats are from CUB-200-2011 [1], Cars-196 [2] and Oxford-IIIT Pet [3] respectively.

TABLE V
PERFORMANCES OF OBJECT-PART SPATIAL CONSTRAINT MODEL, PART ALIGNMENT AND THEIR COMBINATION.

Method	Accuracy (%)		
	CUB-200-2011	Cars-196	Oxford-IIIT Pet
OPSC+PA (ours)	80.65	84.26	85.75
OPSC (ours)	79.74	83.34	83.46
PA (our previous [15])	65.41	68.32	75.42

2) *Effectiveness of object-part spatial constraint model and part alignment*: Compared with our previous conference paper [15] which only performs part alignment for selecting discriminative parts, we further employ object-part spatial constraint model to drive the discriminative parts selection. The object spatial constraint ensures selected parts with high representativeness, while part spatial constraint eliminates redundancy and enhances discrimination of selected parts. Both of them are jointly employed to exploit the subtle and local discrimination for distinguishing the subcategories. In Fig. 9 and Table V, “OPSC” refers to the object-part spatial constraint model, “PA” refers to part alignment which is adopted by our previous conference paper [15], and “OPSC+PA” refers to combining the above two ones, which is adopted by our OPADDL approach. From the left columns of 3 datasets in Fig. 9, we can see that only performing part alignment in part-level attention model, without applying object-part spatial constraint model, causes the selected parts: (1) have large areas of background noise but small areas of object, (2) have large overlap with each other which leads to the redundant information. And from Table V, we can see that the classification accuracies of parts selected by object-part spatial constraint model (“OPSC”) are better than parts selected with part alignment (“PA”) on all 3 datasets. And applying part alignment on the basis of object-part spatial constraints, which means the part-level attention model in our approach, further improves the classification performance. This verifies that aligning the discriminative parts with same semantic meaning

TABLE VI
PERFORMANCES OF PATCH FILTER.

Method	Accuracy (%)		
	CUB-200-2011	Cars-196	Oxford-IIIT Pet
ft-patches	80.82	86.79	88.14
ft-original	80.11	85.76	87.52

together can further improves the results of part-level attention model.

3) *Effectiveness of patch filter*: Through patch filter in the object-level attention model, some image patches are selected from the candidate image patches. These patches are relevant to objects, and provide multiple views and scales of original images. These relevant patches are used to train *ClassNet* to boost the effectiveness of *ClassNet*. In Table VI, “ft-patches” refers to fine-tuning on image patches selected through patch filter in object-level attention model and “ft-original” refers to fine-tuning only on original images. The results are the classification accuracy of prediction on original images. Fine-tuning on the selected image patches achieves better accuracy due to the effectiveness of multi-view and multi-scale feature learning based on the patch filter in our OPADDL approach.

V. CONCLUSION

In this paper, the OPADDL approach has been proposed for weakly supervised fine-grained image classification, which jointly integrates two level attention models: object-level localizes objects of images, and part-level selects discriminative parts of object. The two level attentions jointly improve the multi-view and multi-scale feature learning and enhances their mutual promotion. Besides, the discriminative parts selection is driven by the object-part spatial model, which combines two spatial constraints: object spatial constraint ensures selected parts highly representative, and part spatial constraint eliminates redundancy and enhances discrimination of selected parts. Combining the two spatial constraints promotes the subtle and local discrimination localization. Importantly, our OPADDL avoids the heavy labor consuming of labeling to march toward practical application. Comprehensive experimental results show effectiveness of our OPADDL approach compared with more than 10 state-of-the-art methods on 3 widely-used datasets.

The future work lies in two aspects: First, we will focus on learning better fine-grained representation via more effective and precise parts localization methods. Second, we will also attempt to apply semi-supervised learning into our work to make full use of large amounts of web data. Both of them will be employed to further improve the fine-grained image classification performance.

REFERENCES

- [1] Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The caltech-ucsd birds-200-2011 dataset. 2011.
- [2] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3d object representations for fine-grained categorization. *International Conference of Computer Vision Workshop (ICCV)*, pages 554–561, 2013.
- [3] Omkar M Parkhi, Andrea Vedaldi, Andrew Zisserman, and CV Jawahar. Cats and dogs. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3498–3505, 2012.

- [4] Subhransu Maji, Esa Rahtu, Juho Kannala, Matthew Blaschko, and Andrea Vedaldi. Fine-grained visual classification of aircraft. *arxiv:1306.5151*, 2013.
- [5] Ning Zhang, Jeff Donahue, Ross Girshick, and Trevor Darrell. Part-based r-cnns for fine-grained category detection. *International Conference on Machine Learning (ICML)*, pages 834–849, 2014.
- [6] Xiaopeng Zhang, Hongkai Xiong, Wengang Zhou, Weiyao Lin, and Qi Tian. Picking deep filter responses for fine-grained image recognition. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1134–1142, 2016.
- [7] Xiaopeng Zhang, Hongkai Xiong, Wengang Zhou, and Qi Tian. Fused one-vs-all features with semantic alignments for fine-grained visual categorization. *IEEE Transactions on Image Processing (TIP)*, 25(2):878–892, 2016.
- [8] Jasper RR Uijlings, Koen EA van de Sande, Theo Gevers, and Arnold WM Smeulders. Selective search for object recognition. *International Journal of Computer Vision (IJCV)*, 104(2):154–171, 2013.
- [9] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 580–587, 2014.
- [10] Steve Branson, Grant Van Horn, Serge Belongie, and Pietro Perona. Bird species categorization using pose normalized deep convolutional nets. *arxiv:1406.2952*, 2014.
- [11] Jonathan Krause, Hailin Jin, Jianchao Yang, and Li Fei-Fei. Fine-grained recognition without part annotations. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5546–5555, 2015.
- [12] Feng Zhou and Yuanqing Lin. Fine-grained image classification by exploring bipartite-graph labels. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1124–1133, 2016.
- [13] Yu Zhang, Xiu-Shen Wei, Jianxin Wu, Jianfei Cai, Jiangbo Lu, Viet-Anh Nguyen, and Minh N Do. Weakly supervised fine-grained categorization with part-based image representation. *IEEE Transactions on Image Processing (TIP)*, 25(4):1713–1725, 2016.
- [14] Lingxi Xie, Qi Tian, Richang Hong, Shuicheng Yan, and Bo Zhang. Hierarchical part matching for fine-grained visual categorization. *International Conference of Computer Vision (ICCV)*, pages 1641–1648, 2013.
- [15] Tianjun Xiao, Yichong Xu, Kuiyuan Yang, Jiaxing Zhang, Yuxin Peng, and Zheng Zhang. The application of two-level attention models in deep convolutional neural network for fine-grained image classification. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 842–850, 2015.
- [16] David G Lowe. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision (IJCV)*, 60(2):91–110, 2004.
- [17] Lingxi Xie, Qi Tian, Meng Wang, and Bo Zhang. Spatial pooling of heterogeneous features for image classification. *IEEE Transactions on Image Processing (TIP)*, 23(5):1994–2008, 2014.
- [18] Shenghua Gao, Ivor Wai-Hung Tsang, and Yi Ma. Learning category-specific dictionary and shared dictionary for fine-grained image categorization. *IEEE Transactions on Image Processing (TIP)*, 23(2):623–634, 2014.
- [19] Ning Zhang, Ryan Farrell, Forrest Iandola, and Trevor Darrell. Deformable part descriptors for fine-grained recognition and attribute prediction. *International Conference of Computer Vision (ICCV)*, pages 729–736, 2013.
- [20] Marcel Simon and Erik Rodner. Neural activation constellations: Unsupervised part model discovery with convolutional networks. *International Conference of Computer Vision (ICCV)*, pages 1143–1151, 2015.
- [21] Max Jaderberg, Karen Simonyan, Andrew Zisserman, et al. Spatial transformer networks. *Neural Information Processing Systems (NIPS)*, pages 2017–2025, 2015.
- [22] Shaoli Huang, Zhe Xu, Dacheng Tao, and Ya Zhang. Part-stacked cnn for fine-grained visual categorization. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1173–1182, 2016.
- [23] Han Zhang, Tao Xu, Mohamed Elhoseiny, Xiao lei Huang, Shaoting Zhang, Ahmed Elgammal, and Dimitris Metaxas. Spda-cnn: Unifying semantic part detection and abstraction for fine-grained recognition. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1143–1152, 2016.
- [24] Dequan Wang, Zhiqiang Shen, Jie Shao, Wei Zhang, Xiangyang Xue, and Zheng Zhang. Multiple granularity descriptors for fine-grained categorization. *International Conference on Computer Vision (ICCV)*, pages 2399–2406, 2015.
- [25] Tsung-Yu Lin, Aruni RoyChowdhury, and Subhransu Maji. Bilinear cnn models for fine-grained visual recognition. *International Conference of Computer Vision (ICCV)*, pages 1449–1457, 2015.
- [26] Bo Zhao, Jiashi Feng, Xiao Wu, and Shuicheng Yan. A survey on deep learning-based fine-grained object classification and semantic segmentation. *International Journal of Automation and Computing*, pages 1–17, 2017.
- [27] Pierre Sermanet, Andrea Frome, and Esteban Real. Attention for fine-grained categorization. *arxiv:1412.7054*, 2014.
- [28] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [29] Xiao Liu, Tian Xia, Jiang Wang, and Yuanqing Lin. Fully convolutional attention localization networks: Efficient attention localization for fine-grained recognition. *arxiv:1603.06765*, 2016.
- [30] Bo Zhao, Xiao Wu, Jiashi Feng, Qiang Peng, and Shuicheng Yan. Diversified visual attention networks for fine-grained object classification. *arxiv:1606.08572*, 2016.
- [31] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 248–255, 2009.
- [32] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arxiv:1409.1556*, 2014.
- [33] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arxiv:1502.03167*, 2015.
- [34] Thomas Berg and Peter Belhumeur. Poof: Part-based one-vs.-one features for fine-grained categorization, face verification, and attribute estimation. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 955–962, 2013.
- [35] Luming Zhang, Yang Yang, Meng Wang, Richang Hong, Liqiang Nie, and Xuelong Li. Detecting densely distributed graph patterns for fine-grained image categorization. *IEEE Transactions on Image Processing (TIP)*, 25(2):553–565, 2016.
- [36] Hossein Azizpour and Ivan Laptev. Object detection using strongly-supervised deformable part models. *European conference on computer vision (ECCV)*, pages 836–849, 2012.
- [37] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1–9, 2015.
- [38] Ken Chatfield, Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Return of the devil in the details: Delving deep into convolutional nets. *arxiv:1405.3531*, 2014.
- [39] Hantao Yao, Shiliang Zhang, Yongdong Zhang, Jintao Li, and Qi Tian. Coarse-to-fine description for fine-grained visual categorization. *IEEE Transactions on Image Processing (TIP)*, 25(10):4858–4872, 2016.
- [40] Zhe Xu, Dacheng Tao, Shaoli Huang, and Ya Zhang. Friend or foe: Fine-grained categorization with weak supervision. *IEEE Transactions on Image Processing (TIP)*, 26(1):135–146, 2017.
- [41] Yin Cui, Feng Zhou, Yuanqing Lin, and Serge Belongie. Fine-grained categorization and dataset bootstrapping using deep metric learning with humans in the loop. *arxiv:1512.05227*, 2015.
- [42] Zhe Xu, Shaoli Huang, Ya Zhang, and Dacheng Tao. Webly-supervised fine-grained visual categorization via deep domain adaptation. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2016.
- [43] Di Lin, Xiaoyong Shen, Cewu Lu, and Jiaya Jia. Deep lac: Deep localization, alignment and classification for fine-grained recognition. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1666–1674, 2015.
- [44] Saining Xie, Tianbao Yang, Xiaoyu Wang, and Yuanqing Lin. Hyperclass augmented and regularized deep learning for fine-grained image classification. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2645–2654, 2015.
- [45] Jonathan Krause, Timnit Gebru, Jia Deng, Li-Jia Li, and Li Fei-Fei. Learning features and parts for fine-grained recognition. *International Conference on Pattern Recognition (ICPR)*, pages 26–33, 2014.
- [46] Yaming Wang, Jonghyun Choi, Vlad Morariu, and Larry S Davis. Mining discriminative triplets of patches for fine-grained classification. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1163–1172, 2016.
- [47] Yaming Wang, Vlad I Morariu, and Larry S Davis. Weakly-supervised discriminative patch learning via cnn for fine-grained recognition. *arxiv:1611.09932*, 2016.

- [48] Jinjun Wang, Jianchao Yang, Kai Yu, Fengjun Lv, Thomas Huang, and Yihong Gong. Locality-constrained linear coding for image classification. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3360–3367, 2010.
- [49] Lingxi Xie, Richang Hong, Bo Zhang, and Qi Tian. Image classification and retrieval are one. *ACM on International Conference on Multimedia Retrieval*, pages 3–10, 2015.
- [50] Hossein Azizpour, Ali Sharif Razavian, Josephine Sullivan, Atsuto Maki, and Stefan Carlsson. From generic to specific deep representations for visual recognition. *IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPR)*, pages 36–45, 2015.
- [51] Qi Qian, Rong Jin, Shenghuo Zhu, and Yuanqing Lin. Fine-grained visual categorization via multi-stage metric learning. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3716–3724, 2015.
- [52] Ahmet Iscen, Giorgos Tolias, Philippe-Henri Gosselin, and Hervé Jégou. A comparison of dense region detectors for image search and fine-grained classification. *IEEE Transactions on Image Processing (TIP)*, 24(8):2369–2381, 2015.
- [53] Naila Murray and Florent Perronnin. Generalized max pooling. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2473–2480, 2014.
- [54] Liefeng Bo, Xiaofeng Ren, and Dieter Fox. Multipath sparse coding using hierarchical matching pursuit. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 660–667, 2013.
- [55] Anelia Angelova and Shenghuo Zhu. Efficient object detection and segmentation for fine-grained recognition. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 811–818, 2013.